

TITLE OF THE INVENTION

Root Cause Correlation in Connectionless Networks.

FIELD OF THE INVENTION

5 The present invention relates to computer network technology in general, and in particular to correlation of network errors to root causes in connectionless networks.

BACKGROUND OF THE INVENTION

10 Connectionless computer networks, such as Internet Protocol (IP) networks, are typically formed by connecting multiple routers to each other using either point-to-point connections or the Data Link Layer of the International Standard Organization's Open System Interconnect (ISO/OSI) network model, commonly referred to as "layer 2." One of the main features of a connectionless network is the ability of a network node, such as a PC, to connect directly to any of the routers and send/receive packetized data to/from any other network node connected to any other router. To accomplish this each node is typically uniquely identified by a unique network address, known in IP networks as an IP address.

15 Routing of packets in a connectionless computer network is now described by way of example with reference to Fig. 1. When a node A sends a packet to a node B, A must specify the address of B as the destination address of the packet. The first router R1 that accepts the packet forwards the packet to the next router R2 on the path to B, whereupon R2 forwards the packet to the next router R3 on the path to B, and so on. When the packet reaches the router to which B is directly connected, it is forwarded to B. It may thus be seen that, for any given destination address to which a packet is addressed, every router in the network should know the packet's next "hop," i.e., to which next router the packet is to be forwarded. Each router typically maintains this information in a routing table which contains a mapping between addresses or address groups, such as IP subnets, and the next hop for packets destined for these addresses.

25 When a link connecting two routers in a network fails, a partitioning of the network may occur. Thus in Fig. 1, if the link between R1 and R2 fails, nodes A and C can

still communicate with each other but not with nodes B and D, and vice versa. Each router will typically automatically detect this situation and update its routing table accordingly, such as by eliminating entries whose next hop is unreachable. However, nodes in one partition may still try to send packets to nodes in the other partition. When this occurs, a
5 "no route to destination" error is typically generated and logged by the first router to detect the problem, which then reports the problem to the network management system (NMS). The NMS must then decide what action to take, such as tracing the error to its root cause. In large networks where there may be many active communication sessions between nodes at one time, a single link failure event might cause numerous "no route to destination"
10 notifications to be generated in every router in one partition which receives packets that are destined for the other partition and reported to the NMS. Thus, where the existence of a link failure is already known to the NMS, it would be advantageous to know whether or not a routing error is caused by the link failure, as well as which nodes might be affected by the link failure, obviating the need for the NMS to take action that it would normally take.

15

SUMMARY OF THE INVENTION

The present invention provides for the correlation of routing errors to link failures in a connectionless network.

In one aspect of the present invention a method is provided for correlating
20 routing errors to link failures in a network, the method including detecting a link failure between a first and a second router in a network, associating a first node address indicated in a first routing table of the first router with a first partition of the network, where a next hop of a packet destined for the first node address is the second router, associating a second node address indicated in a second routing table of the second router with a second
25 partition of the network, where a next hop of a packet destined for the second node address is the first router, and correlating an error notification resulting from the failed delivery of a packet with the link failure where a source address of the packet corresponds to the first node address and a destination address of the packet corresponds to the second node address.

30

In another aspect of the present invention any of the steps are performed with respect to a connectionless network.

In another aspect of the present invention a method the correlating step includes correlating a "no route to destination" error.

In another aspect of the present invention the associating steps comprise constructing a connectivity table.

5 In another aspect of the present invention the method further includes suppressing the error.

In another aspect of the present invention any of the steps are performed in a distributed network management system by at least one software agent associated with either of the routers.

10 In another aspect of the present invention the method further includes notifying at least one other agent in the network of the associations of the nodes to the partitions, where the other agent is not associated with either of the routers.

In another aspect of the present invention a method is provided for correlating routing errors to link failures in a network, the method including identifying a path between
15 a first node and a second node in a network, detecting a link failure in the network, determining if the link failure lay along the path, and correlating an error notification resulting from the failed delivery of a packet with the link failure where a source address of the packet corresponds to an address of the first node, where a destination address of the packet corresponds to an address of the second node, and where the link failure lay along
20 the path.

In another aspect of the present invention the identifying step includes identifying either of a most commonly used route and a most heavily used route between the nodes in accordance with a predefined measure of use.

In another aspect of the present invention any of the steps are performed with
25 respect to a connectionless network.

In another aspect of the present invention the correlating step includes correlating a "no route to destination" error.

In another aspect of the present invention the method further includes suppressing the error.

In another aspect of the present invention any of the steps are performed in a distributed network management system by a software agent associated with either of the routers.

5 In another aspect of the present invention a system is provided for correlating routing errors to link failures in a network, the system including means for detecting a link failure between a first and a second router in a network, means for associating a first node address indicated in a first routing table of the first router with a first partition of the network, where a next hop of a packet destined for the first node address is the second router, means for associating a second node address indicated in a second routing table of
10 the second router with a second partition of the network, where a next hop of a packet destined for the second node address is the first router, and means for correlating an error notification resulting from the failed delivery of a packet with the link failure where a source address of the packet corresponds to the first node address and a destination address of the packet corresponds to the second node address.

15 In another aspect of the present invention any of the means are operative with respect to a connectionless network.

In another aspect of the present invention the means for correlating is operative to correlate a "no route to destination" error.

20 In another aspect of the present invention the means for associating are operative to construct a connectivity table.

In another aspect of the present invention the system further includes means for suppressing the error.

25 In another aspect of the present invention a system any of the means are operative in a distributed network management system including at least one software agent associated with either of the routers.

In another aspect of the present invention the system further includes means for notifying at least one other agent in the network of the associations of the nodes to the partitions, where the other agent is not associated with either of the routers.

30 In another aspect of the present invention a system is provided for correlating routing errors to link failures in a network, the system including means for identifying a path between a first node and a second node in a network, means for detecting a link failure in

the network, means for determining if the link failure lay along the path, and means for correlating an error notification resulting from the failed delivery of a packet with the link failure where a source address of the packet corresponds to an address of the first node, where a destination address of the packet corresponds to an address of the second node, and where the link failure lay along the path.

In another aspect of the present invention the means for identifying is operative to identify either of a most commonly used route and a most heavily used route between the nodes in accordance with a predefined measure of use.

In another aspect of the present invention any of the means are operative with respect to a connectionless network.

In another aspect of the present invention the means for correlating step is operative to correlate a "no route to destination" error.

In another aspect of the present invention the system further includes means for suppressing the error.

In another aspect of the present invention any of the means are operative in a distributed network management system including a software agent associated with either of the routers.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully from the following detailed description taken in conjunction with the appended drawings in which:

Fig. 1 is a simplified pictorial illustration of a network framework, useful in understanding present invention;

Fig. 2 is a simplified pictorial illustration of a network framework supporting error correlation, constructed and operative in accordance with a preferred embodiment of the present invention;

Fig. 3 is a simplified flowchart illustration of a method of correlation of routing errors to link failures in a connectionless network, operative in accordance with a preferred embodiment of the present invention.

Fig. 4 is a simplified flowchart illustration of a method of correlation of routing errors to link failures in a connectionless network supported by a distributed network

management system, operative in accordance with a preferred embodiment of the present invention; and

Fig. 5 is a simplified flowchart illustration of a method of identifying nodes that may be affected by link failures in a connectionless network, operative in accordance with a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Reference is now made to Fig. 2, which is a simplified pictorial illustration of a network framework supporting error correlation, constructed and operative in accordance with a preferred embodiment of the present invention, and additionally to Fig. 3, which is a simplified flowchart illustration of a method of correlation of routing errors to link failures in a connectionless network, operative in accordance with a preferred embodiment of the present invention. In Fig. 2 a link 200 between two routers R1 and R2 is shown as having failed, as designated by an 'x' through link 200. Prior to the failure of link 200, a routing table 202 of router R1 shows that the next hop for packets destined for B and D is R2, while a routing table 204 of router R2 shows that the next hop for packets destined for A and C is R1. It may be seen that two partitions 206 and 208 (shown in dashed lines) are thus created in that nodes A and C cannot communicate with nodes B and D via link 200, and vice versa.

A network management system (NMS) 210 preferably maintains copies of routing tables 202 and 204. Having detected a link failure between R1 and R2, NMS 210 may create a connectivity table 212 indicating which nodes are in each of partitions 206 and 208. Since NMS 210 knows that R2 is inaccessible to R1 via link 200, NMS 210 may associate with partition 206 those node addresses in its copy of routing table 202 whose next hop is R2. Likewise, NMS 210 may associate with partition 208 those node addresses in routing table 204 whose next hop is R1. Should NMS 210 receive a "no route to destination" error notification from a network router together with the source and destination addresses of the packet that could not be delivered, NMS 210 may look up the source and destination addresses in connectivity table 212 to determine whether they are from different partitions. If both the source and destination addresses are from different partitions, then the "no route to destination" error notification may be an attempt to send

the packet across failed link 200. Thus, the error notification may be correlated with the link failure that is already known to NMS 210, and the error may be suppressed and need not be investigated further. Alternatively, the error notification should not be correlated with the link failure and may be investigated or otherwise acted upon by NMS 210.

5 Reference is now made to Fig. 4, which is a simplified flowchart illustration of a method of correlation of routing errors to link failures in a connectionless network supported by a distributed network management system, operative in accordance with a preferred embodiment of the present invention. In Fig. 4 the present invention is implemented in a distributed network management system, such as is described in U.S.
10 Patent Application No. 09/799,637 and published as Published Application No. 20010039577, where every router has an associated software agent which continuously monitors the state of the router and its links. The agents monitoring R1 and R2 would thus detect the failure of link 200 and then communicate with each other to create connectivity table 212 which may then be provided to the agents of all other routers in the network.
15 Thus, when any router Rx encounters a "no route to destination" error, its associated agent looks up the source and destination addresses in connectivity table 212 to determine whether they are from different partitions, and action may be taken or the error notification ignored as described above.

 Reference is now made to Fig. 5, which is a simplified flowchart illustration of a
20 method of identifying nodes that may be affected by link failures in a connectionless network, operative in accordance with a preferred embodiment of the present invention. In Fig. 5 a list of virtual paths in a network is maintained, where each virtual path represents the traversal of the links, routers, and other network elements comprising the most commonly used and/or most heavily used routes between network nodes, as determined
25 using any predefined measure of use. The virtual path list may be maintained centrally, such as by NMS 210, or in a distributed manner, such as by one or more agents in a distributed network management system. The virtual path list may be created using any conventional technique, such as by identifying common access patterns in router access lists, analyzing network failure alarms (e.g., packet lost, no route, etc.) to determine traffic flow, and
30 determining network tomography from traffic counter patterns. When a failed link is detected, each virtual path may be checked using any known technique to determine if it is

broken and, if so, which nodes and other network elements along the path are affected. Thereafter, should a "no route to destination" error be encountered where the source address of the packet being sent belongs to the node at one end of a virtual path known to have a failed link, and the packet's destination address belongs to the node at the other end
5 of the virtual path, the error may be correlated to the failed link and action may be taken or suppressed as described hereinabove.

It is appreciated that one or more of the steps of any of the methods described herein may be omitted or carried out in a different order than that shown, without departing from the true spirit and scope of the invention.

10 While the methods and apparatus disclosed herein may or may not have been described with reference to specific hardware or software, it is appreciated that the methods and apparatus described herein may be readily implemented in hardware or software using conventional techniques.

15 While the present invention has been described with reference to one or more specific embodiments, the description is intended to be illustrative of the invention as a whole and is not to be construed as limiting the invention to the embodiments shown. It is appreciated that various modifications may occur to those skilled in the art that, while not specifically shown herein, are nevertheless within the true spirit and scope of the invention.